Decoding long-duration GW from BNS with machine learning: Parameter estimation and equations of state

Qian Hu

University of Glasgow

Qian.Hu@glasgow.ac.uk

BNS challenges in 3G

- BNS signals can last for hours in 3G detectors (starting from 5Hz)
- We are expecting >200k BNS events per year in 3G era
- How challenging?
  - A dedicated ROQ-based parameter estimation (PE) cost 1600 CPU hours (PRL. 127 2021 8, 081102), not including the Earth rotation effects
  - Inferring equation of state (EOS) also involves stochastic sampling, which takes O(1)-O(10) hours
- What is the cost? **Optimistically** assuming 1000 CPU hours to process each event (PE+EOS) and 150W CPU power, the 200k BNS will cost (per analysis run)
  - 200 million CPU hours
  - 30 GWh of electricity
  - 4.8 million USD in electricity charges

Machine learning based pipeline

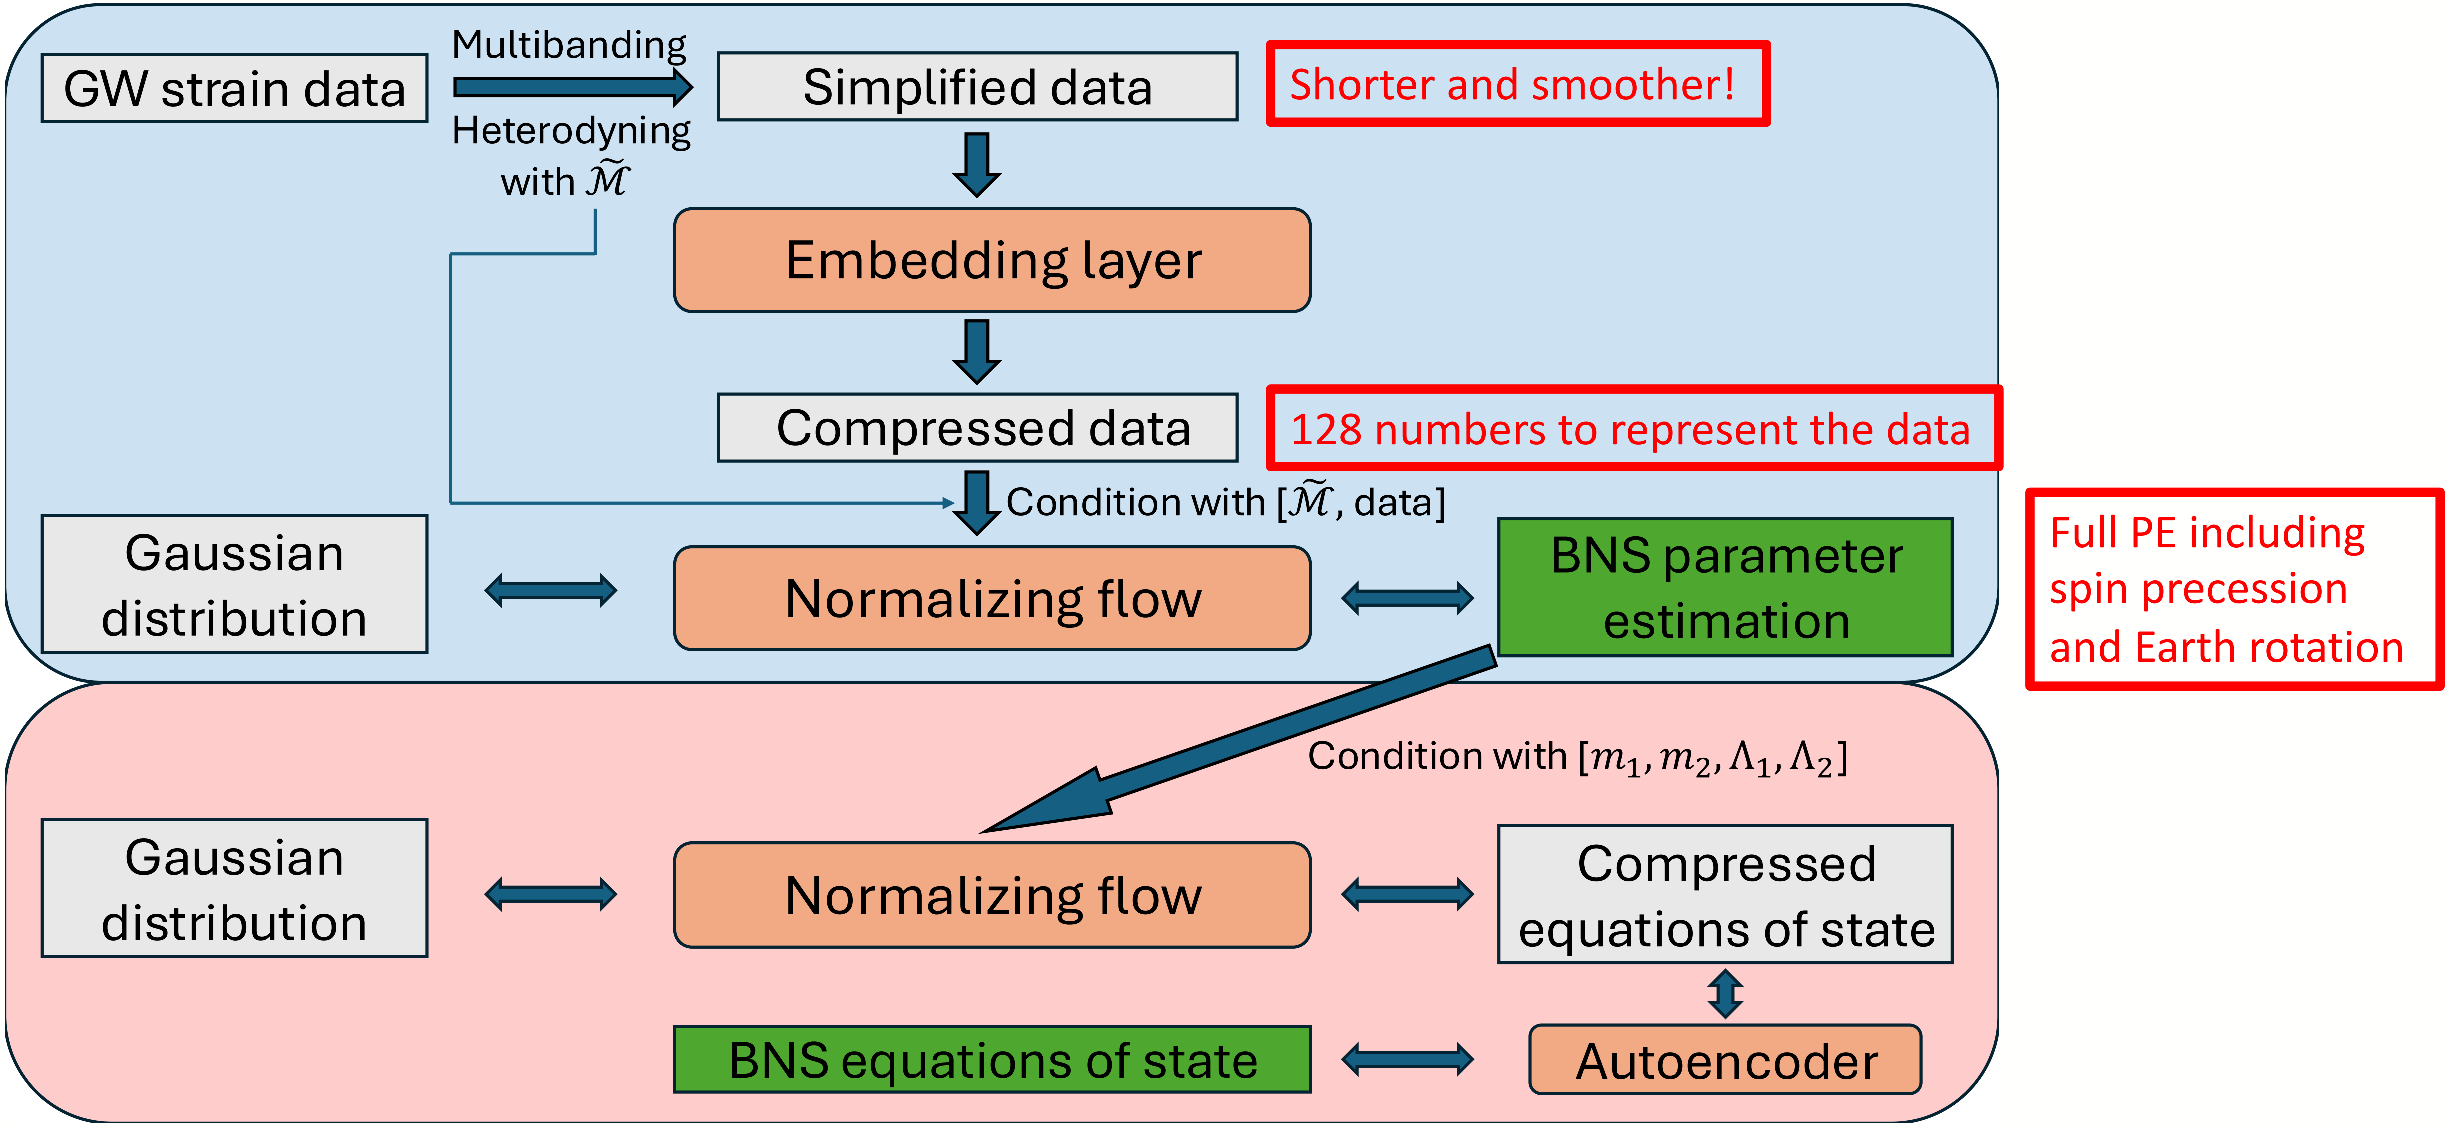

Normalizing flow

- Learn an invertible and differentiable transformation between a target distribution and a Gaussian distribution $p(target) \leftrightarrow p(Gaussian)$

Gaussian

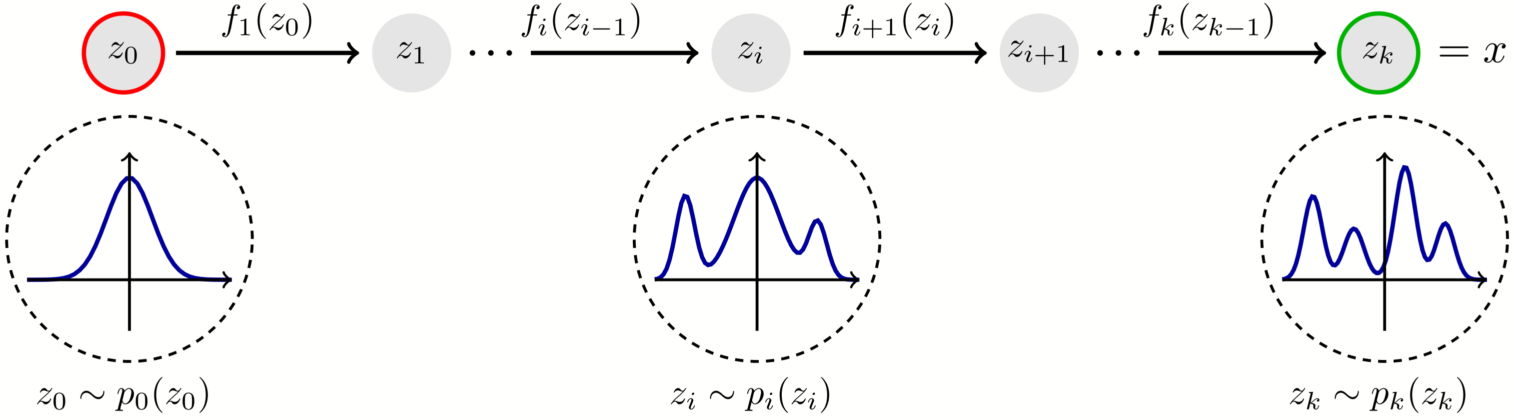

Target distribution,
e.g., GW parameter estimation

- Can be conditioned on data: learning $p(target|data) \leftrightarrow p(Gaussian|data)$

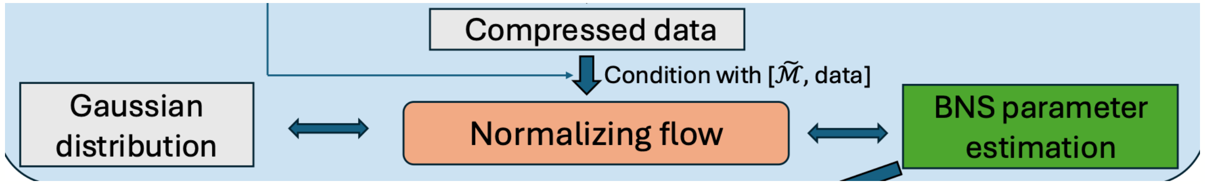

- During inference, samples can be easily drawn from the Gaussian distribution and mapped to data space -> samples of GW parameters

Example parameter estimation

- 1.7+1.6 solar mass at 2500 Mpc, SNR=40
- Takes ~0.3s to generate 5000 samples
- Can constrain source parameters
- Can model degeneracy between parameters
- Crosscheck: models with different embedding layers give consistent results

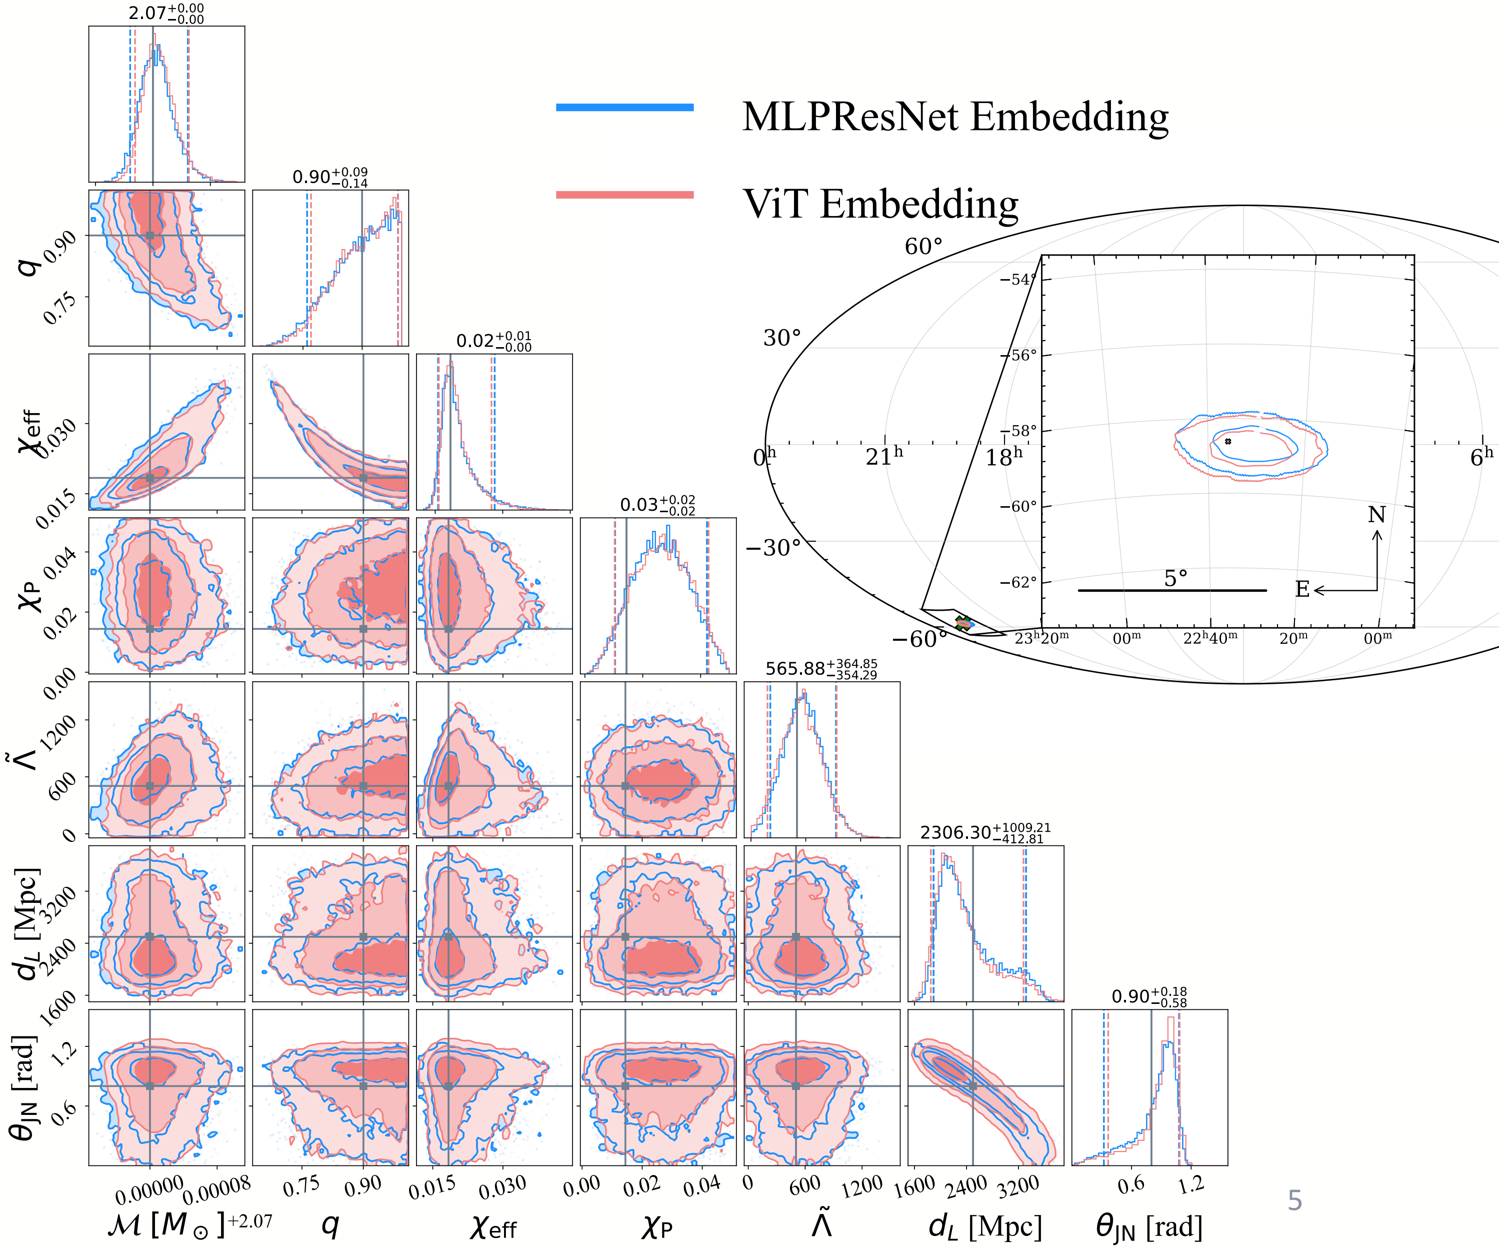

Model validation

- Full PE is prohibitively slow, so we assess our PE models by **precision** and **accuracy**
- Precision: compared with Fisher matrix and SealGW (a fast localization algorithm – check out my poster!)
- Accuracy: p-p plot

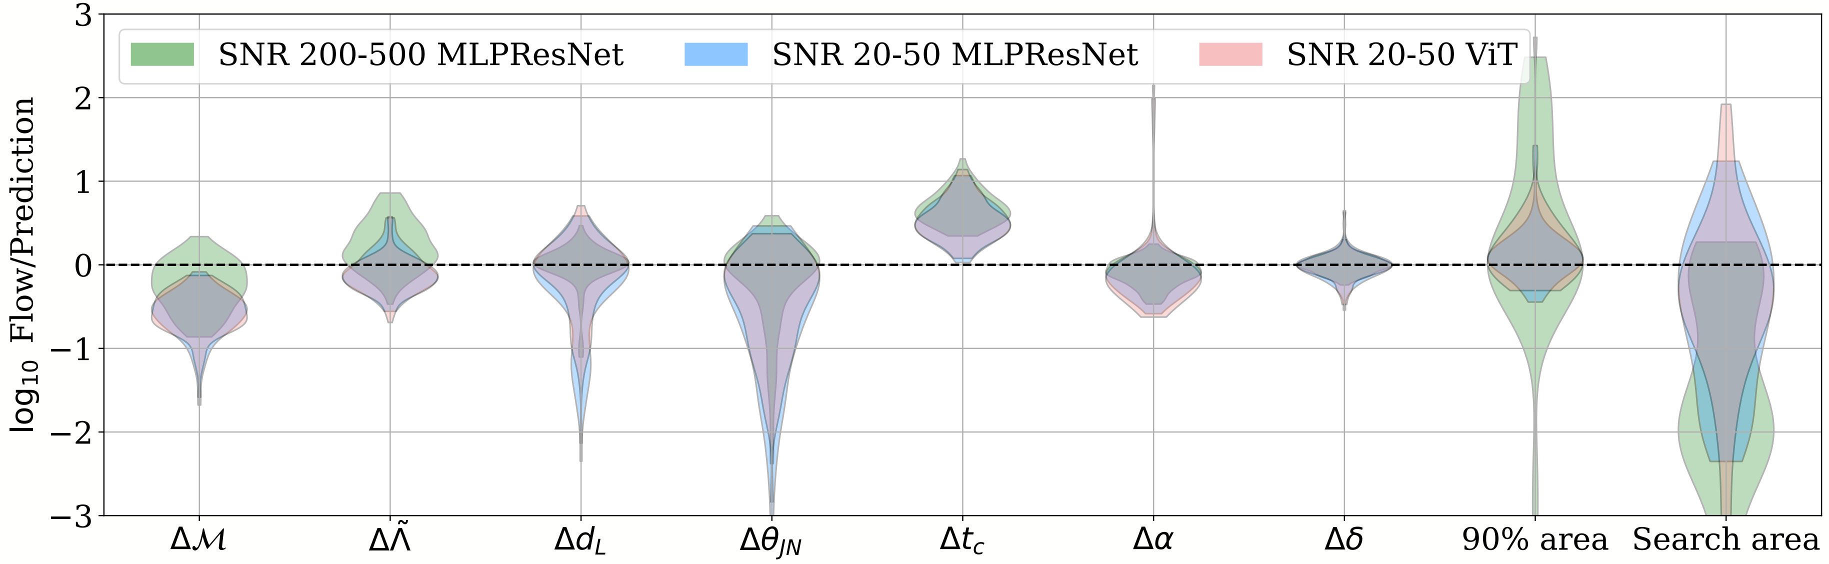

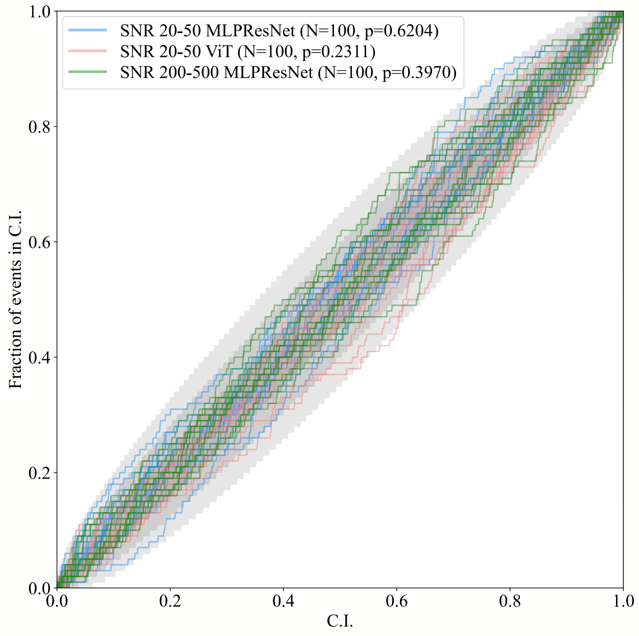

Inferring equations of state

- Posterior samples from PE can be used to infer EOS of neutron stars
- EOS: relation between pressure and density of neutron star matter – without GW it is hard to probe into the dense core!
- We infer the compressed expression of the EOS based on GW PE samples – using normalizing flows!

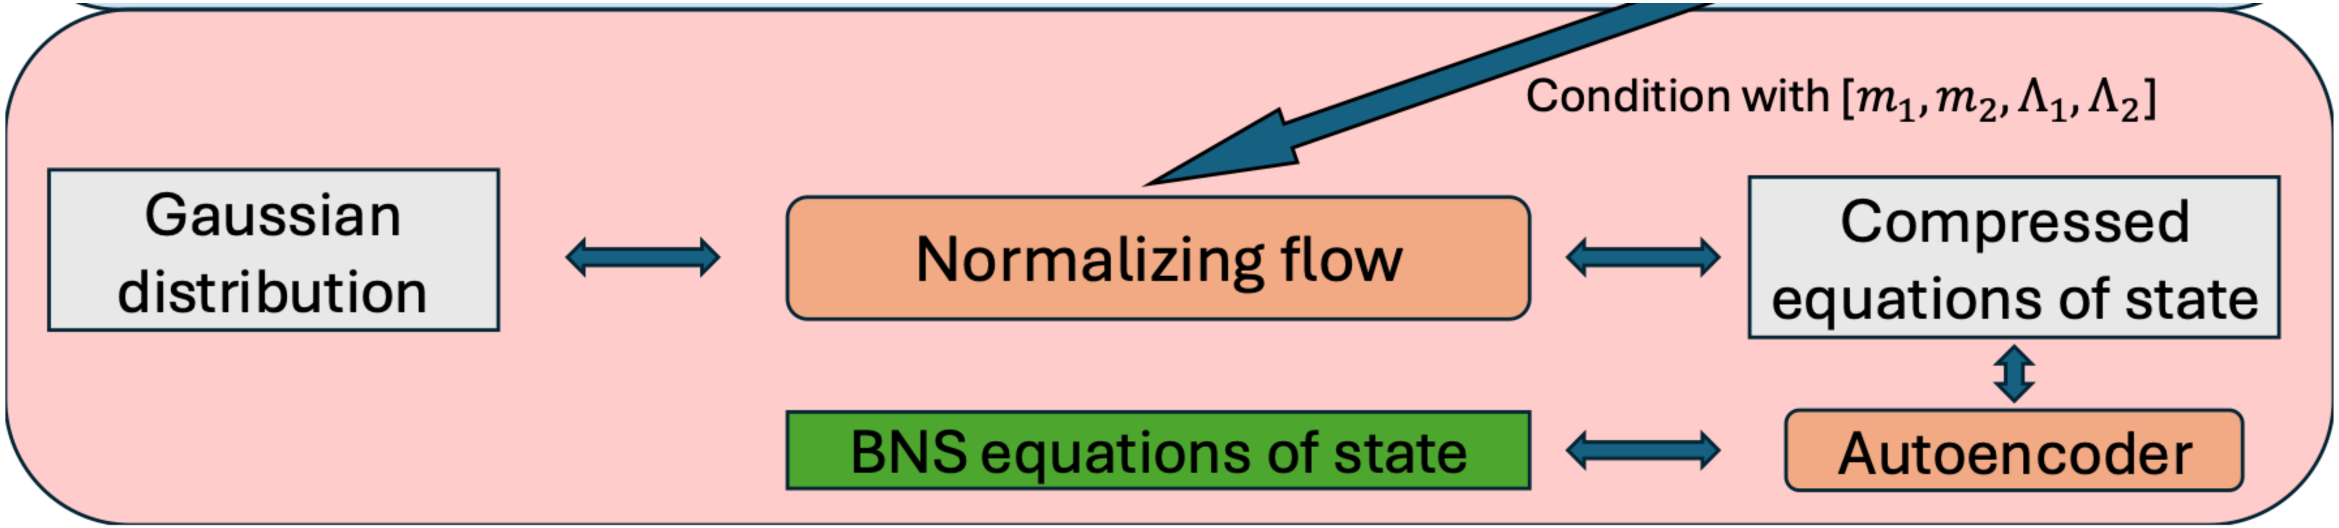

Example EOS constraint

- Simulated two BNS with same underlying EOS but different SNRs: 39 and 390
- We can obtain EOS constraint within 1s!
- High SNR gives tighter constraint as expected

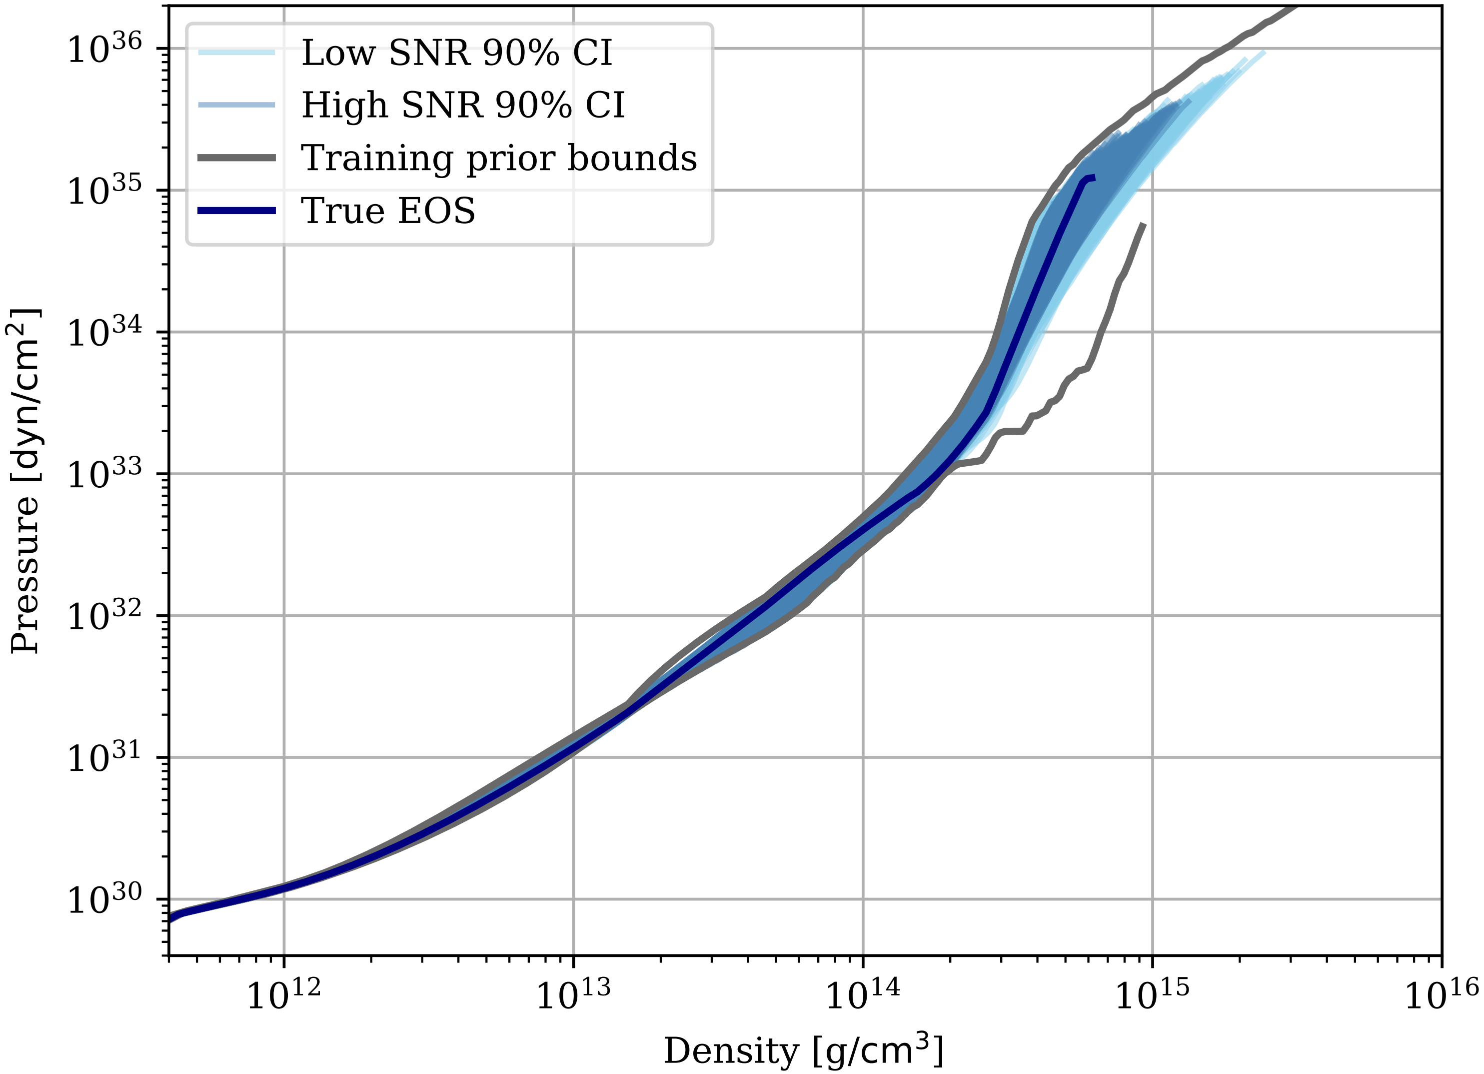

Reduced cost

- Before: **Optimistically** assuming 1000 CPU hours to process each event (PE+EOS) and 150W CPU power, the 200k BNS will cost (per analysis run)
  - 30 GWh of electricity, 4.8 million USD in electricity charges
- What is the cost now? Assuming 1s sampling time and 1min pre- and post-processing CPU time for PE+EOS analysis per event. Assume 500 models are needed to cover the entire parameter space each taking 2 weeks training
  - Inference: 508 kWh, costing approximately 81 USD
  - Training: 25.2 MWh and 4k USD
  - Total: 25.7 MWh and 4.1k USD
  - **Less than 1/1000 of the original cost!**

Summary

- Normalizing flow based analysis pipeline for full PE (precession and earth rotation included) and EOS inference for long BNS signals in 3G
- Validated against Fisher matrix and SealGW because full PE is too expensive
- Energy cost: less than 1/1000 of traditional method for expected 3G catalog

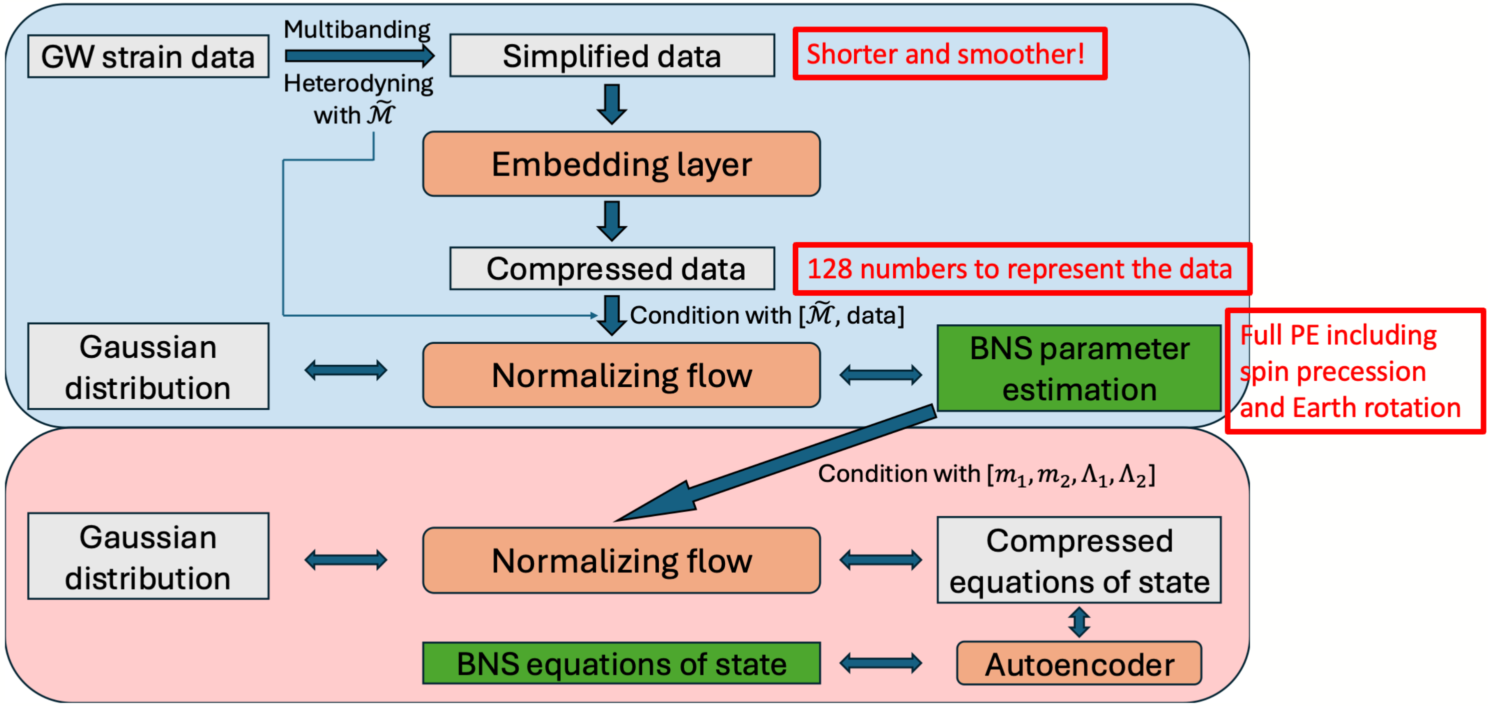

GW data compression

Preprocessing

- 5Hz-1024Hz band, 1.1+1.1Msun –> 12 million data points
- Multibanding:
  - You don’t need high sampling rate in low frequency band
  - Adaptively choosing frequency resolution: 12 million -> 6000
- Heterodyning (relative binning)
  - BNS waveform in frequency domain is highly oscillatory
  - Multiplying the signal with the inverse of dominate oscillation term $e^{i\frac{3}{128}\left(\frac{\pi G M f}{c^3}\right)^{-\frac{5}{3}}}$ to cancel the oscillations

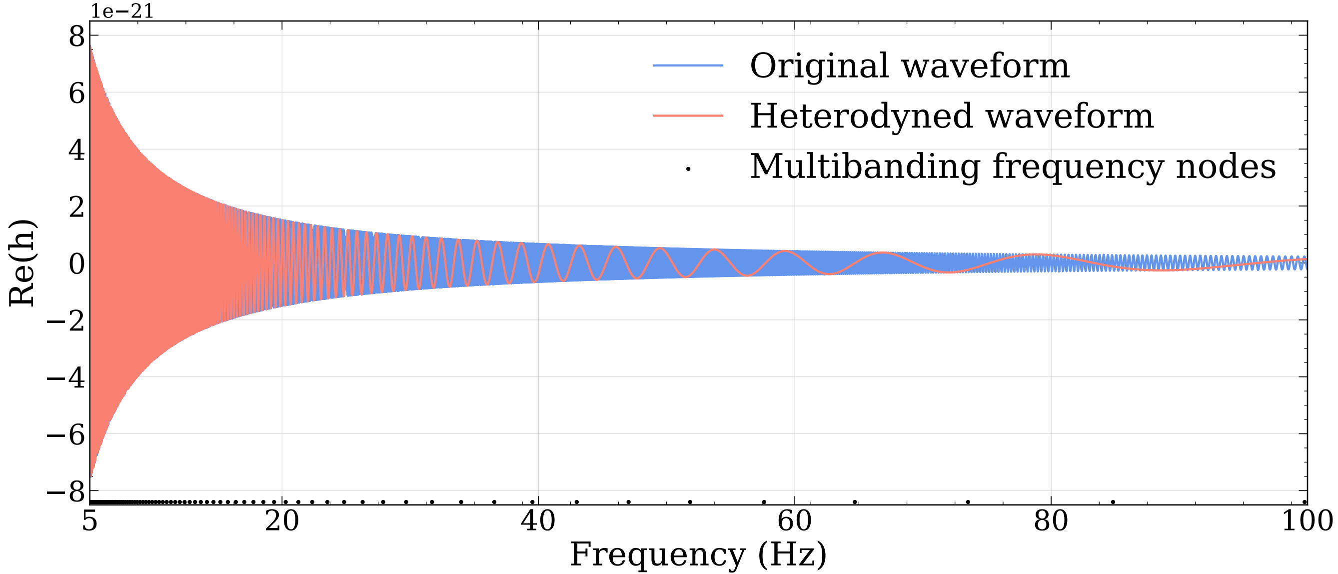

GW data compression

Compression: linear and non-linear

- After multibanding and heterodyning, GW data is short and (relatively) smooth
- We use singular value decomposition (SVD) to extract the linear projections of the data: 6000->128
- We use neural networks to combine different data streams (1 triangular ET + 2 CEs) and compress them: 128*5*2 (data is complex) -> 128 real numbers
  - Residual network of multi-layer perceptron (MLPResNet)
  - Vision Transformer (ViT)

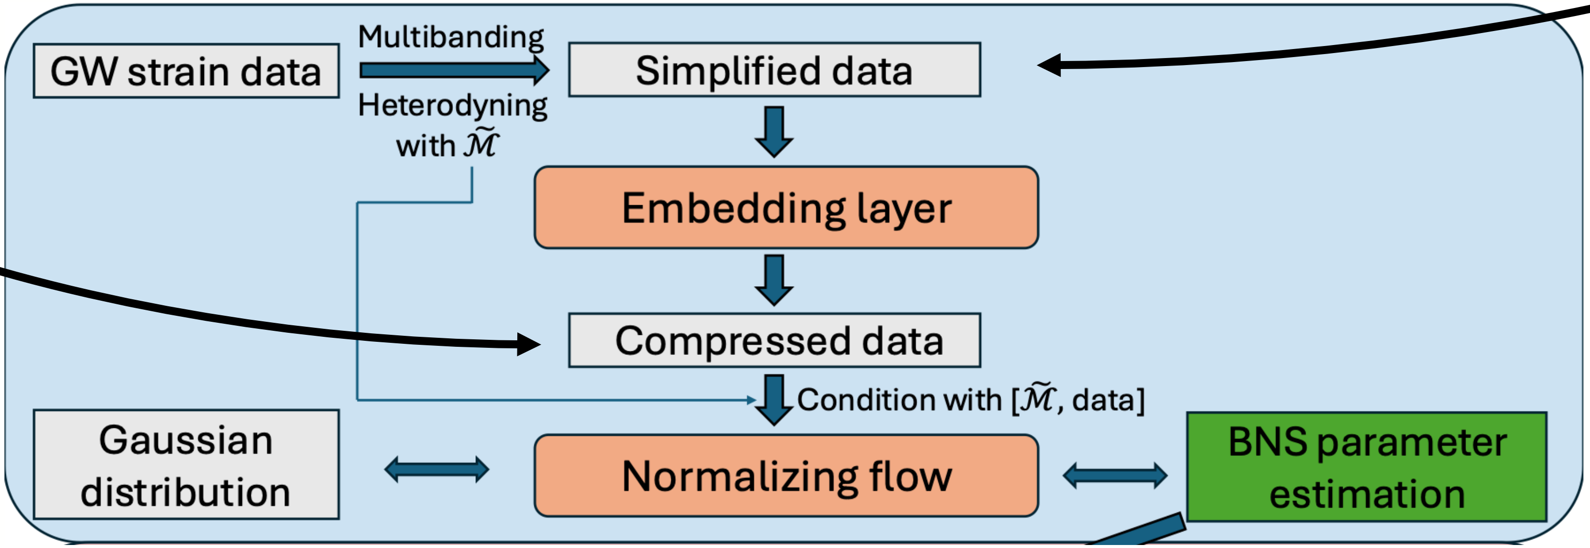

Training

- Training set should be a comprehensive representation of the parameter space, similar to a template bank generation
  - Challenging in high SNR case, low masses, and high dimensions – We meet all three factors!
  - Large training set is required
- We restrict our prior to reduce training set
  - **Low SNR model** (SNR 20-50), chirp mass 2-2.1 solar masses in detector frame, isotropic spin, magnitude < 0.05, random simulating extrinsic parameters during training -> 64 million intrinsic parameters needed
  - **High SNR model** (SNR 200-500), chirp mass 1.3-1.31 solar masses in detector frame, same settings for other parameters -> 100 million intrinsic parameters needed
- Training takes 2-3 weeks on a large GPU

Prior conditioning

- The chirp mass for heterodyning is unknown during inference
- Following DINGO-BNS, we train our model to adapt to small inaccuracies in the chirp mass used for heterodyning
- During inference, we can divide chirp mass space into several segments and perform PE for each. Then choose the one with the maximum likelihood

The Earth rotation effects

- Earth rotation -> changes in response functions of GW detectors
- Encodes information of source location
- For long signals, the Earth rotation's effects need to be considered

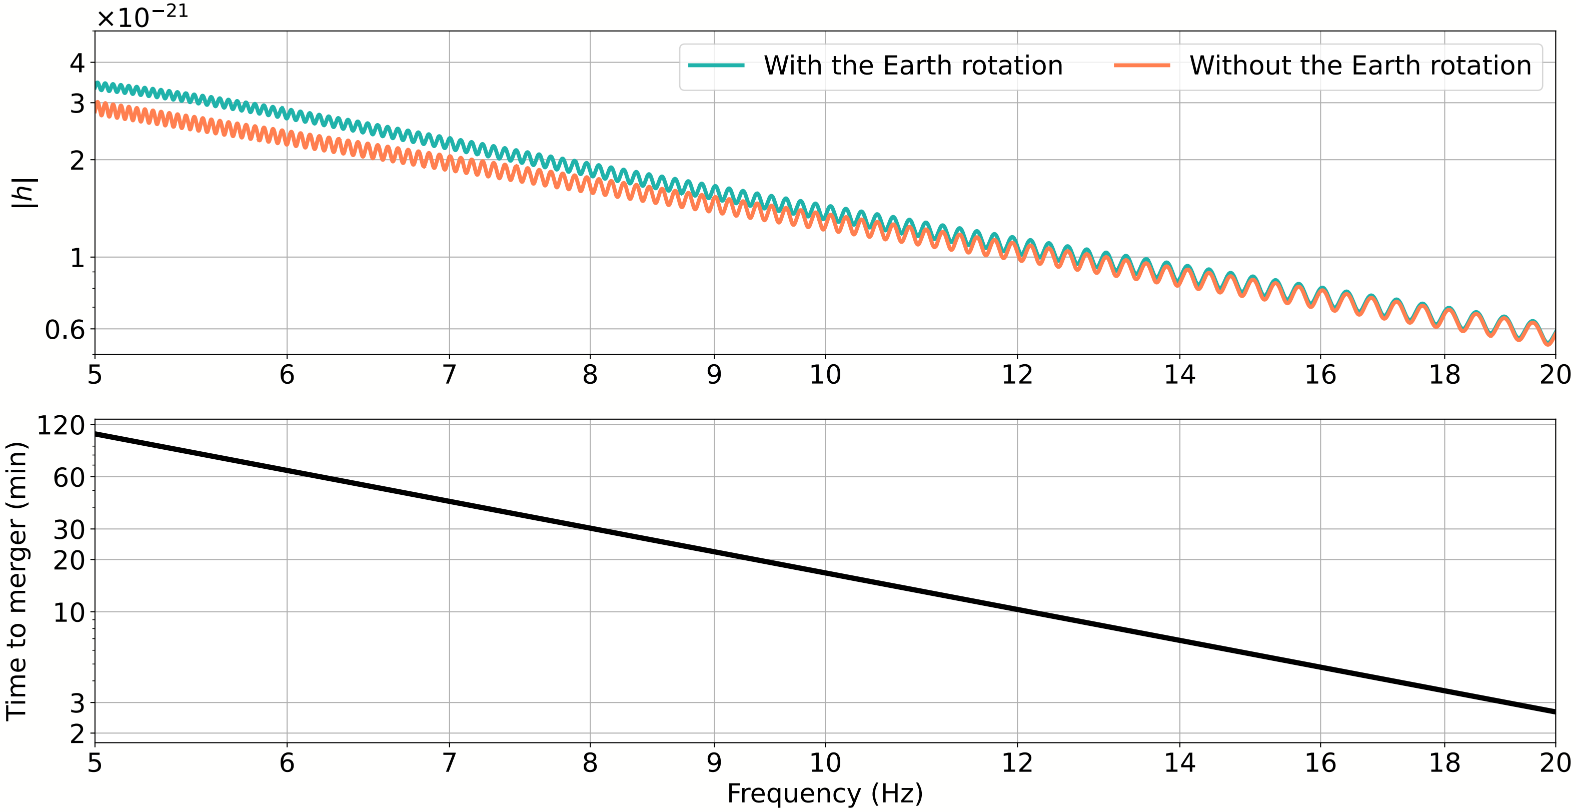

3G PE cost

In prep

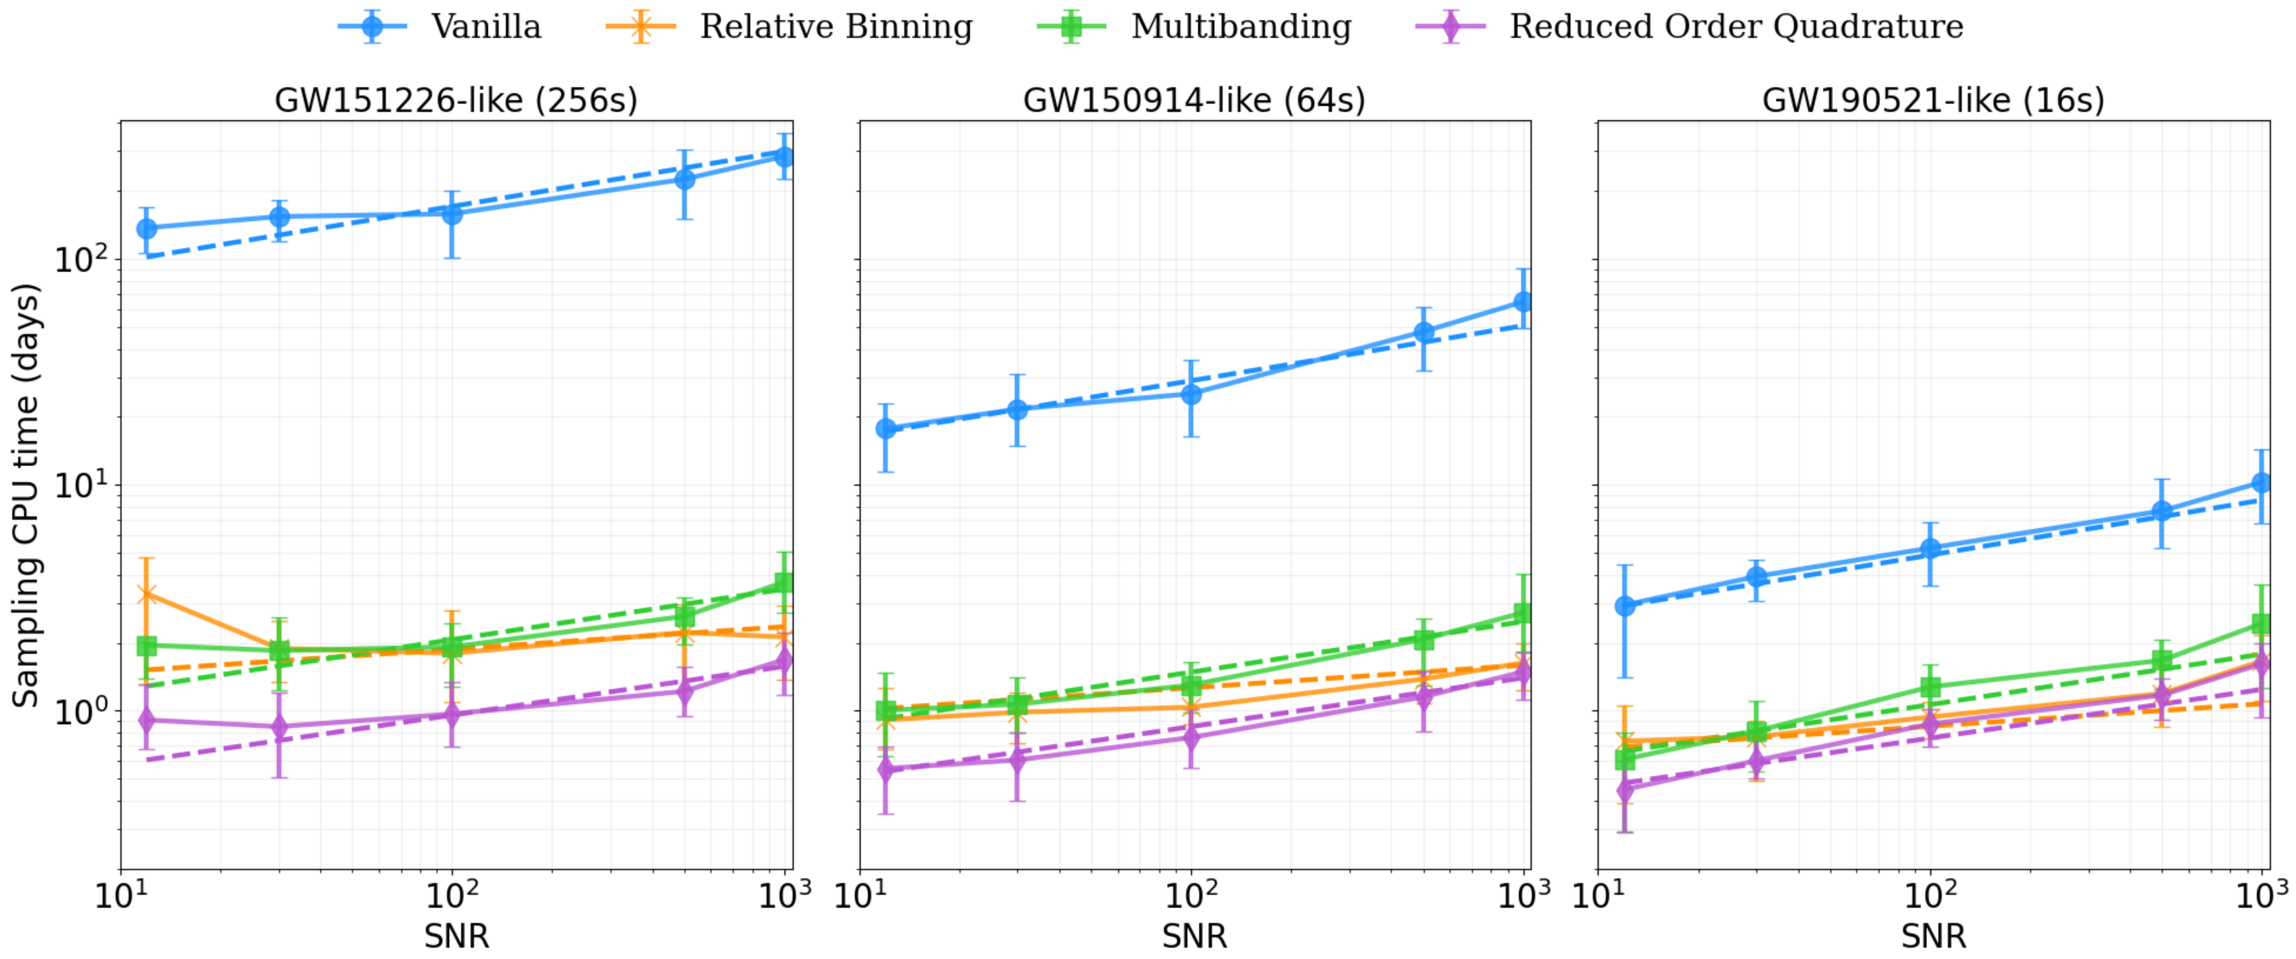

EOS training data

- The EOS for training is sourced from CUTER (Davis+ 2024), which consists of a meta-model and piecewise polytrope structure
- We sample polytrope parameters to generate EOS, for each EOS we solve TOV equation to generate source parameters (masses and Lambdas).
- The source parameters and compressed EOS are used to train the flow model
- During inference, the flow simply takes the source parameters from PE and generate compressed EOS